

spring
Internet-Draft
Intended status: Standards Track
Expires: 11 August 2026

Z. Zhang
Juniper Networks
R. Hoffman
G. Bajwa
TELUS
D. Voyer
Bell Canada
S. Zadok
Broadcom
A. Wang
China Telecom
L. Jalil
Verizon
S. Li
Arrcus
S. Sivabalan
Ciena
7 February 2026

MicroTap Segment in Segment Routing
draft-zzhang-spring-microtap-segment-04

Abstract

This document specifies a microTap segment that can be used to instruct a transit node to make a copy of a segment-routed packet and deliver it to a specified node for the purpose of network monitoring, trouble shooting, or lawful intercept.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 11 August 2026.

Copyright Notice

Copyright (c) 2026 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	3
2. Specification	5
2.1. SR-MPLS Signaling	5
2.1.1. OSPF Signaling	5
2.1.2. ISIS Signaling	7
2.1.3. BGP Signaling	9
2.2. Controller Signaling	10
2.2.1. BGP-LS	10
2.2.2. PCEP	12
2.3. Procedures	12
2.3.1. Optional IOM header	13
2.4. MicroTapping with SRv6	14
2.4.1. ISIS Signaling	14
2.4.2. OSPFv3 Signaling	16
2.4.3. End.TAP (Full SID) Endpoint Behavior	16
2.4.4. End.TAP.X (Full SID) Endpoint Behavior	16
2.4.5. End.TAP (C-SID) Endpoint Behavior	17
2.4.6. End.TAP.X (C-SID) Endpoint Behavior	17
3. Security Considerations	18
4. IANA Assignments	18
5. Contributors	18
6. Appendix A: MicroTap Use Cases	18
6.1. A.1 Use Case 1: Multiple Tapping Nodes with Single Monitor	19
6.1.1. A.1.1 Scenario Description	19
6.1.2. A.1.2 Network Topology	19
6.1.3. A.1.3 Configuration Details	19
6.1.4. A.1.4 Packet Flow Analysis	20
6.2. A.2 Use Case 2: Local Monitor with End.TAP.X	20
6.2.1. A.2.1 Scenario Description	20
6.2.2. A.2.2 Network Topology	20
6.2.3. A.2.3 Configuration Details	21

6.2.4.	A.2.4 Packet Flow Analysis	21
6.3.	A.3 Use Case 3: Single Tapping Node with Multiple Monitors	21
6.3.1.	A.3.1 Scenario Description	21
6.3.2.	A.3.2 Network Topology	21
6.3.3.	A.3.3 Configuration Details	21
6.3.4.	A.3.4 Packet Flow Analysis	22
6.4.	A.4 Use Case Combinations and Real-World Deployments . .	23
6.4.1.	A.4.1 Combining Multiple Use Cases	23
7.	References	23
7.1.	Normative References	23
7.2.	Informative References	24
	Authors' Addresses	24


1. Introduction

Network operators may for various reasons benefit from the ability to tap packets at strategic locations within their respective networks. Segment routing [RFC8402] technology offers the ability to both simplify and improve the operational experience of deploying targeted packet tapping.

The tapping can be only for some random packets for monitoring purposes, so we use the term microTap and tap interchangeably in this document.

The introduction and strategic placement within a SID-list of one or more microTap SIDs can signal the desire to tap traffic at targeted points within the network without the need for explicit configuration on those nodes.

Consider an SR network in the following example diagram where traffic is steered along some paths by using a SID-list in the packets. For network debugging/monitoring purposes, the operator may at any time want for a certain node (e.g., R2 or R3) in the network to tap a copy of a packet to a monitor (e.g. connected to R6), while continue to forward the original packet along its path to the destination.



To make it very flexible and precise on specifying which packets to tap on what node and avoid the need to configure filters on the microTap node, a microTap SID can be inserted to the SID-list after a Node-SID (for the microTap node) or an Adjacency-SID (that leads to the microTap node). When the microTap SID becomes the current active SID, the node does the following:

- * Replicate the packet, and send the copy to the remote monitor
- * Pop the microTap SID off the original packet and continue forwarding

There could be multiple monitors. A microTap SID is associated with a particular monitor (vs. a microTap node). In the above example, there could be another monitor attached to R5. In that case, there would be two microTap SIDs - one for the monitor attached at R5 (say microTap SID S5) and one for the monitor attached at R6 (say microTap SID S6).

The monitor could be a separate server attached to an interface on R5 or R6, or could be an internal service entity on R5 or R6 (which can be viewed as connected via an internal interface). How that is done is outside the scope of this document.

If S5 becomes the active SID in a packet arriving at R2, R2 will tap the packet to R5, by imposing R5's node SID label on top of S5. When the tapped copy arrives at R5, R5 knows that the packet should be sent to the internal or external monitor (because S5, which R5 advertises, becomes the active SID). Similarly, if S6 becomes the active SID in a packet arriving at R3, R3 will tap the packet to R6, by imposing R6's node SID label on top of S6. In case of SRv6, a separate IPv6 header is used to send the packet to the router to which the monitor is attached.

It is possible that a monitor node itself may be on the path of a packet and need to do a tap. This is referred to as local tapping and a separate local microTapping SID is needed - a packet with that microTapping SID active is tapped to a local monitor. For comparison, when a packet was tapped by another node and the tapped copy arrives on the monitor node, it is simply forwarded but not tapped to the monitor.

A microTap SID is advertised by the router that hosts the monitor. It should only become the active SID in a packet arriving at the desired microTap node or the advertising/owning node. A node supporting microTap functionality advertises its ability to do so, so that incapable nodes will never see a microTap SID as the active SID in a packet.

The SID-list may contain multiple microTap SIDs that may or may not be adjacent in the list. For nonadjacent microTap SIDs, different nodes will tap to the same or different monitors (depending on the value of microTap SIDs). For adjacent microTap SIDs in the list, they are likely for different monitors - for the "continue forwarding" part of the first microTap SID, the second microTap SID becomes active segment, leading to the second microTap operation.

2. Specification

2.1. SR-MPLS Signaling

A node (e.g. R2/R3) supporting microTap function advertise its capability to other nodes.

A node (e.g. R5/R6) hosting a monitor is provisioned with a microTap SID allocated from the SRGB. The microTap SID is advertised to other nodes.

A microTap SID MUST be associated with only one specific monitor.

If the same microTap SID value is advertised by more than one node, it MUST be treated by a receiving node as an error and ignored, and MUST NOT be used in the SID-List of a packet.


SRv6 related signaling details will be added in future revisions.

2.1.1. OSPF Signaling

This document defines a new TLV for the advertisement of a microTap SID (from a node hosting a monitor) and an existing TLV is leveraged for the advertisement of tapping capability (from a microTap node).

2.1.1.1. MicroTap-SID TLV

The microTap SID is advertised in a newly defined MicroTap-SID Sub-TLV that mimics the Prefix SID Sub-TLV as defined in Section 5 of [RFC8665]:



where:

Type: To be assigned by IANA

Length: 7 or 8 octets depending on the size of SID (see below).

Flags: Single-octet field. Currently no flags are defined.

Reserved: SHOULD be set to 0 on transmission and MUST be ignored on reception

MT-ID: Multi-Topology ID (as defined in [RFC4915])

Algorithm: Single octet identifying the algorithm the Prefix-SID is associated with as defined in Section 3.1

A router receiving a Prefix-SID from a remote node and with an algorithm value that the remote node has not advertised in the SR-Algorithm TLV (Section 3.1) MUST ignore the Prefix-SID Sub-TLV.

SID/Index/Label: Currently a 4-octet index defining the offset in the Segment Routing Global Block (SRGB) advertised by this router. In the future the flags field may change the definition of this definition of this field.

The MicroTap-SID Sub-TLV MAY appear where a Prefix-SID Sub-TLV is included to advertises a node SID.

2.1.1.2. MicroTap Capability

A new flag T in the Flags field of the Prefix/Adjacency-SID Sub-TLV indicates that a MicroTap SID is allowed to follow the prefix/adjacency SID in a packet:

0	1	2	3	4	5	6	7
+---+---+---+---+---+---+---+---+							
							T
+---+---+---+---+---+---+---+---+							

2.1.2. ISIS Signaling

ISIS signaling is similar to OSPF, as specified in the following sections.

2.1.2.1. MicroTap-SID

The microTap SID is advertised in a newly defined MicroTap-SID Sub-TLV that mimics the Prefix SID Sub-TLV as defined in Section 2.1 of [RFC8667]:

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1								
Type										Length										Flags										Algorithm									
SID/Index/Label (variable)																																							

where:

Type: To be assigned by IANA.

Length: 5 or 6 depending on the size of the SID (described below)

Flags: 1-octet field. Currently no flags are defined.

Algorithm: the router may use various algorithms when calculating reachability to other nodes or to prefixes attached to these nodes. Algorithm identifiers are defined in Section 3.2. Examples of these algorithms are metric-based Shortest Path First (SPF), various sorts of Constrained SPF, etc. The Algorithm field of the Prefix-SID contains the identifier of the algorithm the router uses to compute the reachability of the prefix to which the Prefix-SID is associated.

At origination, the Prefix-SID Algorithm field MUST be set to 0 or any value advertised in the SR-Algorithm sub-TLV.

A router receiving a Prefix-SID from a remote node and with an algorithm value that such remote node has not advertised in the SR-Algorithm sub-TLV MUST ignore the Prefix-SID sub-TLV.

SID/Index/Label: : Currently a 4-octet index defining the offset in the Segment Routing Global Block (SRGB) advertised by his router. In the future the flags field may change the definition of this definition of this field.

The MicroTap-SID Sub-TLV MAY appear where a Prefix-SID Sub-TLV is included to advertises a node SID.

2.1.2.2. Tapping Capability

Similar to OSPF, a new flag T in the Flags field of the Prefix/Adjacency-SID Sub-TLV indicates that a MicroTap SID is allowed to follow the prefix/adjacency SID in a packet:


```

      0  1  2  3  4  5  6  7
+---+---+---+---+---+---+---+
|   |   |   |   |   |   |T |
+---+---+---+---+---+---+---+

```

2.1.3. BGP Signaling

2.1.3.1. MicroTap-SID

A new MicroTap-SID TLV is defined to advertise a microTap SID. It has the same encoding as the Label-Index TLV except with a different type. The following is copied verbatim from [RFC8669]:

```

      0                               1                               2                               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|          Type          |          Length          |  RESERVED  |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|          Flags          |          Label Index          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|          Label Index          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

where:

Type: To be assigned by IANA.

Length: 7, the total length in octets of the value portion of the TLV.

RESERVED: 8-bit field. It MUST be clear on transmission and MUST be ignored on reception.

Flags: 16 bits of flags. None are defined by this document. The Flags field MUST be clear on transmission and MUST be ignored on reception.

Label Index: 32-bit value representing the index value in the SRGB space.

A MicroTap-SID TLV MAY be included in the BGP Prefix-SID attribute.

2.1.3.2. Tapping Capability

A 'T' flag is defined for the existing Originator SRGB TLV's Flags field to indicate that the originator supports microTapping functionality. Exact bit position for the flag is to be assigned by IANA and registered in the "BGP Prefix-SID Originator SRGB TLV Flags" registry.

2.2. Controller Signaling

A controller needs to know about the nodes (e.g. R2/R3) that support tapping function, and the nodes (e.g. R5/R6) hosting a monitor & relevant microTap SID. This information is advertised to the controller by the link-state routing protocols (ISIS and OSPF) or BGP-LS. The signaling for OSPF and ISIS has been covered in the previous sections of this document. This section covers signaling for BGP-LS and PCEP.

2.2.1. BGP-LS

This document defines a new TLV for the advertisement of a microTap SID (from a node hosting a monitor) and an existing TLV is leveraged for the advertisement of tapping capability (from a microTap node).

2.2.1.1. MicroTap SID

The microTap SID is advertised in a newly defined MicroTap-SID TLV that mimics the Prefix SID TLV as defined in Section 2.3.1 of [RFC9085]:

```

0          1          2          3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                         |
|          Type                         |          Length          |
+-----+-----+-----+-----+-----+-----+-----+-----+
|          Flags                       |          Algorithm          |          Reserved          |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                         |
|                                         |          SID/Index/Label (variable)          |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Where:

Type: To be assigned by IANA

Length: Variable. 7 or 8 octets depending on the label or index encoding of the SID.

Flags: 1-octet value that should be set as:

- * IS-IS MicroTap-SID flags as defined in Section 2.1.2.1.
- * OSPFv2 MicroTap-SID flags as defined in Section 2.1.1.1.
- * OSPFv3 MicroTap-SID flags as defined in Section 2.1.1.1.

Algorithm: 1-octet value identifies the algorithm. The semantics of the algorithm are described in Section 3.1.1 of {{RFC8402}}.

Reserved: 2 octets that MUST be set to 0 and ignored on receipt.

SID/Index/Label:

IS-IS: Label or index value as defined in Section 2.1.2.1.

OSPFv2: Label or index value as defined in Section 2.1.1.1.

OSPFv3: Label or index value as defined in Section 2.1.1.1.

The Flags and, as an extension, the SID/Index/Label fields of this TLV are interpreted according to the respective underlying IS-IS, OSPFv2, or OSPFv3 protocol. The Protocol-ID of the BGP-LS Prefix NLRI is used to determine the underlying protocol specification for parsing these fields.

The MicroTap-SID TLV MAY appear where a Prefix-SID TLV advertises a node SID.

2.2.1.2. Tapping Capability

The Flags of Prefix/Adjacency-SID TLV are interpreted according to the respective underlying IGP specification. The new flag T in the Flags field of the Prefix/Adjacency-SID TLV indicates that a MicroTap SID is allowed to follow the prefix/adjacency SID in a packet.

2.2.2. PCEP

An SR-TE path consists of one or more SIDs and may contain one or more microTap SIDs. The SR-TE path information is exchanged between the PCE and PCC in ERO and RRO subobjects. The SR-ERO subobject and SR-RRO subobject defined in [RFC8664] are used to carry a SID which can be a microTap SID.

2.3. Procedures

The node hosting a monitor treats a microTap SID that it advertises as an adjacency SID. In other words, it sets up its forwarding state for the microTap SID such that packets with the microTap SID as current active SID will be sent to the monitor (after popping the microTap SID). It is the responsibility of the monitor to parse the packet (including the remaining SID-list).

A node supporting microTap functionality sets up its forwarding state for each microTap SID that it receives, such that packets with the microTap SID as current active SID are processed as following:

- * Make a copy and send it to the advertising node of the microTap SID. In case of SR-MPLS, this is done by imposing the advertising node's node SID (optionally after imposing the node SID of the microTap node so that the monitor knows the microTap node). In case of SRv6, this is done by imposing an outer IPv6 encapsulation with the destination address being the advertising node's address.
- * Forward the original packet after popping the microTap SID

If a node does not support microtapping but does recognize the microtap SID signaling, the forwarding behavior for the SID is simply pop on that node. This is to safeguard the situation in case the node received a packet with the active SID being a microtap SID.

The ingress node may add microTap SIDs to the SID-list of a packet based on its monitoring/debugging needs or based on SR policies programmed from a controller.

A microTap SID MUST not be placed in the SID-list after a node or adjacency SID that is for or leads to a node that does not advertise microTap capability. Otherwise, the packet with that SID-list will be discarded by the node.

In case of SRv6, the microTap SID and its preceding node SID MAY be merged into a single IPv6 address in SRH: the locator part identifies the microTap SID and the function part is the 3-octet or 4-octet microTap SID.

2.3.1. Optional IOM header

As replicated packets traverse the network from the microtap node to the monitor nodes, packet loss, packet reordering and buffering can occur. To allow packet analysis equipment that receives these replicated packets to accurately analyze the replicated packet flow, additional information is needed in the replicated packet header to recreate the original conditions of the flow.

RFC9197] defines a header with data fields well suited for this purpose. IOAM includes timestamp data, indicating the arrival time the replicated packet was received at the microtap node. This timestamp can be used to reproduce accurate inter-packet gaps during packet analysis. IOAM also includes a sequence number, indicating the order of replicated packets received by the microtap node. This sequence number can be used by the packet analysis equipment to reorder packets, remove duplicated packets, and to alarm on the condition that replicated packets were lost in transit.

The microTap node MAY include an IOAM header in the replicated packet with following fields:

- * Timestamp Seconds
- * Timestamp Fraction
- * 64-bit sequence number

It is RECOMMENDED that all nodes that perform microtap packet replication be Time of Day (ToD) synchronized via Precision Time Protocol (PTP) for the most accurate recreation of packet conditions during analysis.

The added IOAM header is Edge-to-Edge Option-Type, and in addition to possible IOAM header already present when the packet arrives at the microtap node. In case of MPLS, the added IOAM header is an MPLS extension header [I-D.song-mpls-extension-header] that follows the Node SID of the node that originated the microtap SID. The extension

header is followed by the original label stack and its OUL field (Original Upper Layer protocol type) MUST be set to MPLS. In other words, there may be two label stacks in the packet arriving at the node hosting the monitoring station.

If MTU is a concern, the original label stack (except the microTap SID) and extension headers MAY be removed.

2.4. MicroTapping with SRv6


[RFC9800] introduced the concepts of Global Identifiers Block (GIB) for the pool of Compressed SID (C-SID) values available for global allocation and Local Identifiers Block (LIB) for the pool of C-SID values available for local allocation.

This document extends the GIB/LIB concepts to traditional non-compressed full SRv6 SID as well. A Global ID can be allocated from the the GIB and used as C-SID or to construct non-compressed SRv6 SIDs.


Each monitor node MUST advertise a GIB ID for other microTapping nodes to tap traffic to this monitor node. If the monitor node may be on the normal forwarding path of packets and need to tap locally, it MUST also advertises a LIB ID. This GIB ID or LIB ID is referred to as the Tapping ID (TID), and is used by relevant nodes to construct an uncompressed full microTapping SID or as a C-SID for microTapping purposes, as explained below.

2.4.1. ISIS Signaling

A new flag bit, the T-flag, is defined for the Flags field of the locator entry in the SRv6 Locator TLV [RFC9352]:



Flags: 1 octet. The following flags are defined:



D-flag: "up/down bit" as described in Section 4.1 of [RFC5305].

T-flag: Indicates the advertising node supports microTapping for this locator (hence this locator can be used together with the GIB TIDs advertised by monitor nodes).

For the SRv6 End SID sub-TLV, two new Endpoint Behaviors are defined for global microTapping End.TAP (by a tapping node to a monitor via a monitor node) or local microTapping End.TAP.X (by a monitor node to its local monitor) respectively. Each is advertised by the monitor node with a 128-bit End SID in the LB:LN:FUNC:ARG format where the FUNC bits encode the GIB/LIB TID respectively. An SRv6 SID Structure sub-sub-TLV MUST be included to identify the FUNC bits.

The monitor node M installs an IPv6 route LB:LN:GIB-TID::/prefixLength with the End.TAP endpoint behavior for a monitor node and an IPv6 route LB:LN:LIB-TID::/prefixLength with the End.TAP.X endpoint behavior for a monitor node.

When a node N supporting microTapping receives an End.TAP SID advertised by a monitor node M, for each locator that N advertises with the T-flag, N installs an IPv6 route locator:TID::/prefixLength with the End.TAP behavior for a microTapping node, where the prefixLength is (LBL+LNL+FL). . In the uncompressed SID case, the locator:TID:: is a SID that an ingress node can insert to instruct N to microTap to M. In the C-SID case, the TID is a C-SID that can follow any locator C-SID that is advertised with the T-flag.

When a node N receives an End.TAP.X SID advertised by a monitor node M, it does not install any corresponding forwarding state, except that it notes that the LIB TID in the End.TAP.X SID can be used as C-SID or to construct non-compressed SRv6 SIDs that can be inserted in to the SID list of a packet to instruct M to do local microTapping.

If the End.TAP or End.TAP.X SID is advertised with FlexAlgo 0, it MAY be used together with locators of any FlexAlgo if the locator has the T-flag set. Otherwise, it MUST only be used together with locators of corresponding FlexAlgo that has the T-flag set.

2.4.2. OSPFv3 Signaling

To be added.

2.4.3. End.TAP (Full SID) Endpoint Behavior

If a node N receives a packet whose IPv6 DA D matches a route with the End.TAP endpoint behavior, N does:

```
S01. If (IPv6 Hop Limit <= 1) {
S02.     Send an ICMP Time Exceeded message to the Source Address,
        Code 0 (Hop limit exceeded in transit),
        interrupt packet processing, and discard the packet.
S03. }
S04. Decrement IPv6 Hop Limit by 1.
S05. If N is a microTapping node:
        Replicate the packet.
S06.     For the replicated packet,
        Push the microTapping End SID from the monitor node M
        to the packet using H.encap behavior specified in
        {{RFC8986}}
        Submit the packet to the egress IPv6 FIB lookup for
        transmission to the new destination
S07.     For the original packet,
        Decrement Segment Left
        Copy the Segment Left SID to outer IPv6 DA
        Submit the packet to the egress IPv6 FIB lookup for
        transmission to the new destination
S08. If N is the monitor node:
        Pop the SRH that was inserted by the microTapping node.
        Transmit the packet to the monitor associated with
        the End.TAP SID.
```

2.4.4. End.TAP.X (Full SID) Endpoint Behavior

This is for a monitor node to tap a packet to its local monitor.

If a node N receives a packet whose IPv6 DA D matches a route with the End.TAP.X endpoint behavior, N does:


```
S01. If (IPv6 Hop Limit <= 1) {
S02.   Send an ICMP Time Exceeded message to the Source Address,
      Code 0 (Hop limit exceeded in transit),
      interrupt packet processing, and discard the packet.
S03. }
S04. Decrement IPv6 Hop Limit by 1.
      Replicate the packet.
S05. For the replicated packet,
      Transmit it to the monitor associated with
      the End.TAP.X SID.
S06. For the original packet,
      Decrement Segment Left
      Copy the Segment Left SID to outer IPv6 DA
      Submit the packet to the egress IPv6 FIB lookup for
      transmission to the new destination
```

2.4.5. End.TAP (C-SID) Endpoint Behavior

When the microTap node N receives a packet whose IPv6 DA matches a route with the End.TAP endpoint behavior, N does:

```
S01. If (IPv6 Hop Limit <= 1) {
S02.   Send an ICMP Time Exceeded message to the Source Address,
      Code 0 (Hop limit exceeded in transit),
      interrupt packet processing, and discard the packet.
S03. }
S04. Decrement IPv6 Hop Limit by 1.
S05. If N is a microTapping node:
      Replicate the packet.
S06.   For the replicated packet,
      Push the microTapping End SID from the monitor node M
      to the packet using H.encap behavior specified in
      {{RFC8986}}
      Submit the packet to the egress IPv6 FIB lookup for
      transmission to the new destination
S07.   For the original packet,
      Shift the IPv6 DA for two C-SIDs
      Submit the packet to the egress IPv6 FIB lookup for
      transmission to the new destination
S08. If N is the monitoring node:
      Shift the IPv6 DA for two C-SIDs
      Transmit the packet to the monitor associated with
      the End.TAP C-SID.
```

2.4.6. End.TAP.X (C-SID) Endpoint Behavior

This is for a monitor node to tap a packet to its local monitor.

If a node N receives a packet whose IPv6 DA D matches a route with the End.TAP.X endpoint behavior, N does:

```
S01. If (IPv6 Hop Limit <= 1) {  
S02.   Send an ICMP Time Exceeded message to the Source Address,  
      Code 0 (Hop limit exceeded in transit),  
      interrupt packet processing, and discard the packet.  
S03. }  
S04. Decrement IPv6 Hop Limit by 1.  
      Replicate the packet,  
S05. For the replicated packet,  
      Transmit it to the monitor associated with  
      the End.TAP.X SID.  
      For the original packet,  
      Shift the IPv6 DA for two C-SIDs  
      Submit the packet to the egress IPv6 FIB lookup for  
      transmission to the new destination
```

3. Security Considerations

To be added.

4. IANA Assignments

To be added.

5. Contributors

The following also contributed to this document:

Peter Van Oene
Juniper Networks
pvanoene@juniper.net

Abhishek Chakraborty
Juniper Networks
cabhi@juniper.net

6. Appendix A: MicroTap Use Cases

This appendix provides example use cases for MicroTap with SRv6 uSID using the F3216 format (32-bit uSID Block + 16-bit uSID). In this format, the Locator length L=48, leaving 16 bits for the first uSID, and up to 6×16 -bit uSIDs per container (one IPv6 DA).

***Normative Note on SID Structure Advertisement:** Per RFC 9352 §9 (IS-IS), RFC 9513 §10 (OSPFv3), and RFC 9514 §8 (BGP-LS), every SRv6 SID advertisement SHOULD include the SID Structure sub-TLV (or sub-sub-TLV) to indicate the exact split of LB/LN/FUNC/ARG bits. For domains using F3216, the values are:

- * LB Length = 32 bits
- * LN Length = 16 bits
- * Function Length = 16 bits
- * Argument Length = 0 bits

This ensures that receivers parse each 16-bit uSID correctly as one CSID, and that addresses like ...:200:050c:300:050c:4:: are unambiguous.

6.1. A.1 Use Case 1: Multiple Tapping Nodes with Single Monitor

6.1.1. A.1.1 Scenario Description

In this use case, multiple tapping nodes (R2 and R3) are configured to tap traffic to the same monitor node (R5/Monitor-1) using the same Tapping ID (TID). This scenario is useful when network operators want to collect traffic samples from multiple strategic points in the network for centralized analysis.

6.1.2. A.1.2 Network Topology

```
Src ---- R1 ---- R2 ---- R3 ---- R4 ---- Dst | | | | R5 ---- R6
Monitor-1
```

6.1.3. A.1.3 Configuration Details

***Monitor Node (R5) Configuration:**

- Allocates Tapping ID (TID): 0x050c
- Advertises End.TAP behavior with GIB-TID 0x050c (Include the ***SRv6 SID Structure*** so sources can compress this SID; see RFC 9800 §9.1/§10.)
- Installs route: 2001:cafe:500:050c::/64 with End.TAP behavior

***Tapping Nodes (R2 & R3) Configuration:**

- Both nodes advertise locators with T-flag set
- R2: 2001:cafe:200::/48 with T-flag
- R3: 2001:cafe:300::/48 with T-flag
- Both nodes install routes for the same TID:
- R2: 2001:cafe:200:050c::/64 with End.TAP behavior
- R3: 2001:cafe:300:050c::/64 with End.TAP behavior

6.1.4. A.1.4 Packet Flow Analysis

Initial Packet from R1: IPv6 Header: SA: 2001::1 DA: <LB>;:200:050c:300:050c:4:: Payload: IPv4 SA=A.A.A.A, DA=B.B.B.B

At R2 (First Tapping Node): 1. R2 processes the End.TAP C-SID (050c) 2. Duplicates the packet for monitoring 3. For the replicated packet: - Performs H.encap behavior (RFC 8986) - Creates new IPv6 header with monitor destination - Encapsulates original packet with full context - Sends to monitor: DA=<LB>;:500:050c 4. For the original packet: - Advance to the next 16-bit uSID per RFC 9800 NEXT-uSID procedure (advance to next uSID) - Continues forwarding: DA becomes <LB>;:300:060c:4::

At R3 (Second Tapping Node): 1. R3 processes the End.TAP C-SID (060c) 2. Duplicates the packet for monitoring 3. For the replicated packet: - Performs H.encap behavior (RFC 8986) - Creates new IPv6 header with monitor destination - Encapsulates original packet with full context - Sends to monitor: DA=<LB>;:500:050c 4. For the original packet: - Advance to the next 16-bit uSID per RFC 9800 NEXT-uSID procedure (advance to next uSID) - Continues forwarding: DA=<LB>;:4::

At Monitor Node (R5): - Receives tapped packets from both R2 and R3 - Processes End.TAP behavior, the outer IPv6 header is removed, leaving the original inner IPv6/SRH stack intact, and forwards the original packet (or metadata) to the local monitor.

6.2. A.2 Use Case 2: Local Monitor with End.TAP.X

6.2.1. A.2.1 Scenario Description

In this use case, a node (R2) performs local tapping using the End.TAP.X behavior. The tapped traffic is sent directly to a local monitor without traversing the network to a remote monitoring node. This is ideal for real-time analysis or when network bandwidth to remote monitors is limited.

6.2.2. A.2.2 Network Topology

Src ---- R1 ---- R2 ---- R3 ---- R4 ---- Dst | Local Monitor

6.2.3. A.2.3 Configuration Details

Local Monitor Node (R2) Configuration: - Allocates Local Tapping ID (TID): 0x0d (from LIB) - Advertises End.TAP.X behavior with LIB-TID 0x0d - Locator: 2001:cafe:200::/48 with T-flag - Local monitor address: 2001:1:1:1::1 - Installs route: 2001:cafe:200:0d::/64 with End.TAP.X behavior

6.2.4. A.2.4 Packet Flow Analysis

Initial Packet from R1: IPv6 Header: SA: 2001::1 DA: <LB>::200:0d:4:: Payload: IPv4 SA=A.A.A.A, DA=B.B.B.B

At R2 (Local Tapping and Monitor Node): 1. R2 processes the End.TAP.X C-SID (0x0d) 2. Duplicates the packet for local monitoring 3. For the replicated packet: - Transmits directly to local monitor (no H.Insert) - No network traversal required 4. For the original packet: - Advance to the next 16-bit uSID per RFC 9800 NEXT-uSID procedure (advance to next uSID) - Continues forwarding: DA=<LB>:4::

6.3. A.3 Use Case 3: Single Tapping Node with Multiple Monitors

6.3.1. A.3.1 Scenario Description

In this use case, a single tapping node (R2) is configured to tap traffic to multiple monitor nodes (R5 and R6) simultaneously. This scenario is particularly useful when different types of analysis are required or when traffic needs to be distributed across multiple monitoring systems for load balancing or specialized processing.

6.3.2. A.3.2 Network Topology

```
Src ---- R1 ---- R2 ---- R3 ---- R4 ---- Dst | |---- R5 ----
Monitor-1 | |---- R6 ---- Monitor-2
```

6.3.3. A.3.3 Configuration Details

Monitor Node R5 Configuration: - Allocates Tapping ID (TID): 0x050c - Advertises End.TAP behavior with GIB-TID 0x050c (Include the *SRv6 SID Structure* so sources can compress this SID; see RFC 9800 §9.1/§10.) - Locator: 2001:cafe:500::/48 - Installs route: 2001:cafe:500:050c::/64 with End.TAP behavior - Monitor address: 2001:1:1:1::1

***Monitor Node R6 Configuration:** - Allocates Tapping ID (TID): 0x060c
- Advertises End.TAP behavior with GIB-TID 0x060c (Include the ***SRv6 SID Structure*** so sources can compress this SID; see RFC 9800 §9.1/§10.) - Locator: 2001:cafe:600::/48 - Installs route: 2001:cafe:600:060c::/64 with End.TAP behavior - Monitor address: 2001:1:1:2::1

***Tapping Node R2 Configuration:** - Advertises locator with T-flag set: 2001:cafe:200::/48 - Installs routes for both TIDs: -
2001:cafe:200:050c::/64 with End.TAP behavior (for R5) -
2001:cafe:200:060c::/64 with End.TAP behavior (for R6)

6.3.4. A.3.4 Packet Flow Analysis

Because the two MicroTap CSIDs are ***adjacent***, R2 will execute ***two*** replicate-and-forward actions before advancing the original.

***Initial Packet from R1:** IPv6 Header: SA: 2001::1 DA: <LB>:200:050c:060c:4:: Payload: IPv4 SA=A.A.A.A, DA=B.B.B.B

***At R2 (Tapping Node with Multiple Monitors):** 1. R2 processes the first End.TAP C-SID (050c) 2. Creates first duplicated packet (Dup-P1) for R5: - Performs H.encap behavior (RFC 8986) with (500:050c) - Creates new IPv6 header with monitor destination - Encapsulates original packet with full context - Sends to R5: DA=<LB>:500:050c

1. Creates second duplicated packet (Dup-P2) for R6:

- * Performs H.encap behavior (RFC 8986) with (600:060c)
- * Creates new IPv6 header with monitor destination
- * Encapsulates original packet with full context
- * Sends to R6: DA=<LB>:600:060c

2. For the original packet:

- * Advance to the next 16-bit uSID per RFC 9800 NEXT-uSID procedure (advance to next uSID)
- * Continues forwarding: DA=<LB>:4::

***At R5 (Monitor Node 1):** - Receives tapped packet from R2 - Processes End.TAP behavior, removes the outer encapsulation (H.encap), and forwards the original packet (or metadata) to the local monitoring process at 2001:1:1:1::1

At R6 (Monitor Node 2): - Receives tapped packet from R2 - Processes End.TAP behavior, removes the outer encapsulation (H.encap), and forwards the original packet (or metadata) to the local monitoring process at 2001:1:1:2::1

6.4. A.4 Use Case Combinations and Real-World Deployments

6.4.1. A.4.1 Combining Multiple Use Cases

The three use cases described above can be combined in real-world network deployments to create comprehensive monitoring solutions:

Hybrid Deployment Example: Src ---- R1 ---- R2 ---- R3 ---- R4 ----
 Dst | | | |---- R6 ---- Monitor-2 | |---- R5 ---- Monitor-1 | Local
 Monitor

In this scenario: - R2 performs local tapping (Use Case 2) to its local monitor - R2 also taps to R5 (Use Case 3 - multiple monitors) - R3 taps to R6 (Use Case 1 - single monitor) - R5 can receive traffic from multiple sources (Use Case 1)

7. References

7.1. Normative References

- [RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.
- [RFC8665] Psenak, P., Ed., Previdi, S., Ed., Filsfils, C., Gredler, H., Shakir, R., Henderickx, W., and J. Tantsura, "OSPF Extensions for Segment Routing", RFC 8665, DOI 10.17487/RFC8665, December 2019, <<https://www.rfc-editor.org/info/rfc8665>>.
- [RFC8667] Previdi, S., Ed., Ginsberg, L., Ed., Filsfils, C., Bashandy, A., Gredler, H., and B. Decraene, "IS-IS Extensions for Segment Routing", RFC 8667, DOI 10.17487/RFC8667, December 2019, <<https://www.rfc-editor.org/info/rfc8667>>.
- [RFC8669] Previdi, S., Filsfils, C., Lindem, A., Ed., Sreekantiah, A., and H. Gredler, "Segment Routing Prefix Segment Identifier Extensions for BGP", RFC 8669, DOI 10.17487/RFC8669, December 2019, <<https://www.rfc-editor.org/info/rfc8669>>.

- [RFC9085] Previdi, S., Talaulikar, K., Ed., Filsfils, C., Gredler, H., and M. Chen, "Border Gateway Protocol - Link State (BGP-LS) Extensions for Segment Routing", RFC 9085, DOI 10.17487/RFC9085, August 2021, <<https://www.rfc-editor.org/info/rfc9085>>.
- [RFC9800] Cheng, W., Ed., Filsfils, C., Li, Z., Decraene, B., and F. Clad, Ed., "Compressed SRv6 Segment List Encoding", RFC 9800, DOI 10.17487/RFC9800, June 2025, <<https://www.rfc-editor.org/info/rfc9800>>.
- [RFC9352] Psenak, P., Ed., Filsfils, C., Bashandy, A., Decraene, B., and Z. Hu, "IS-IS Extensions to Support Segment Routing over the IPv6 Data Plane", RFC 9352, DOI 10.17487/RFC9352, February 2023, <<https://www.rfc-editor.org/info/rfc9352>>.
- [RFC8986] Filsfils, C., Ed., Camarillo, P., Ed., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "Segment Routing over IPv6 (SRv6) Network Programming", RFC 8986, DOI 10.17487/RFC8986, February 2021, <<https://www.rfc-editor.org/info/rfc8986>>.

7.2. Informative References

- [RFC8664] Sivabalan, S., Filsfils, C., Tantsura, J., Henderickx, W., and J. Hardwick, "Path Computation Element Communication Protocol (PCEP) Extensions for Segment Routing", RFC 8664, DOI 10.17487/RFC8664, December 2019, <<https://www.rfc-editor.org/info/rfc8664>>.
- [I-D.song-mpls-extension-header]
Song, H., Zhou, T., Andersson, L., Zhang, Z. J., and R. Gandhi, "MPLS Network Actions using Post-Stack Extension Headers", Work in Progress, Internet-Draft, draft-song-mpls-extension-header-13, 11 October 2023, <<https://datatracker.ietf.org/doc/html/draft-song-mpls-extension-header-13>>.

Authors' Addresses

Zhaohui Zhang
Juniper Networks
Email: zzhang@juniper.net

Ryan Hoffman
TELUS
Email: ryan.hoffman@telus.com

Gurminderjit Bajwa
TELUS
Email: gurminderjit.bajwa@telus.com

Daniel Voyer
Bell Canada
Email: daniel.voyer@bell.ca

Shay Zadok
Broadcom
Email: shay.zadok@broadcom.com

Aijun Wang
China Telecom
Email: wangaj3@chinatelecom.cn

Luay Jalil
Verizon
Email: luay.jalil@verizon.com

Shengtao Li
Arrcus
Email: shen@arrcus.com

Siva Sivabalan
Ciena
Email: ssivabal@ciena.com